

200 Gbps TBIRD CubeSat Downlink: Pre-Flight Test Results

Curt M. Schieler, Kathleen M. Riesing, Andrew J. Horvath, Bryan C. Bilyeu, Jesse S. Chang,
Ajay S. Garg, Jade P. Wang, Bryan S. Robinson

Massachusetts Institute of Technology, Lincoln Laboratory
244 Wood Street, Lexington, MA, USA 02421

ABSTRACT

The Terabyte Infrared Delivery (TBIRD) program will establish an optical communication link from a 6U nanosatellite in low-Earth orbit to a ground station at burst rates up to 200 Gbps. The system is capable of reliable data delivery from a 2 TB storage buffer on the payload to a ground terminal in the presence of atmospheric fading. An overview of the communication architecture for TBIRD is provided as well as results from communications performance testing of the 3U lasercom payload prior to spacecraft integration. Launch is scheduled for mid-year 2022.

Keywords: free-space optical communications, low-Earth orbit, coherent modem, ARQ, cubesat

1. INTRODUCTION

Recent years have seen significant growth in space-based laser communication system demonstrations and deployments.¹ The narrow beams and available spectrum of lasercom systems as compared to traditional radio-frequency systems enables smaller, low cost terminals that can operate over long link distances at very high data rates. Various near-Earth applications have been pursued, such as GEO-based relays servicing user missions in low-Earth orbit (LEO)^{2,3} and proliferated LEO constellations that use inter-satellite lasercom crosslinks to enable broadband internet access for ground users.⁴ In both of these examples, multiple network hops are utilized to transfer data from source to destination with low latency.

In contrast, direct-to-Earth (DTE) architectures for data delivery involve a direct space-to-ground link from the spacecraft to one or more ground stations during the course of its orbit. DTE systems in LEO typically have infrequent and short duration contacts with ground stations, which means that sensor data must be buffered on-board the spacecraft between contacts and end users on the ground must be able to tolerate delays between data collection and data reception. However, LEO-DTE systems have the significant advantage that the link distances are relatively small, enabling systems with some combination of higher data rates, lower SWaP space and ground terminals, and relaxed pointing requirements. LEO-DTE lasercom demonstrations in recent years have focused on the benefit of lower SWaP and have miniaturized terminals with very low mass and power consumption and volumes as small as 1U, while still achieving data rates that are competitive with or exceed existing RF solutions (100-1000 Mbps).⁵⁻⁸ Many of these demonstrations are targeting applications for small-sat or CubeSat scale platforms.

NASA's TBIRD program has been developing a LEO-DTE architecture that taps into the extremely high data rates achievable at optical frequencies by leveraging terrestrial fiber telecom technology to achieve a 200-Gbps downlink in a small form-factor terminal.⁹⁻¹² Even though contact times from LEO may be only a few minutes long, bursting data down at rates of this magnitude enables data volumes of multiple terabytes (TB) in a single pass. Key elements of the architecture include 100 Gbps commercial off-the-shelf (COTS) fiber transceivers, a terabyte-class on-board storage buffer capable of high speed readout, and an automatic repeat request (ARQ) protocol that provides error-free data transmission in the presence of atmospheric fading.

To demonstrate this architecture, a TBIRD flight mission was established as part of NASA's Pathfinder Technology Demonstrator (PTD) program.¹³ As part of the PTD-3 mission, the TBIRD lasercom payload,

This material is based upon work supported by the National Aeronautics and Space Administration under Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Aeronautics and Space Administration .

roughly 3U in volume, will be hosted on a 6U CubeSat that is scheduled to launch in June 2022. The ground terminal will be located at JPL’s Optical Communications Telescope Laboratory (OCTL).

In this paper, we give an overview of the payload system architecture and present results from testing of the flight unit prior to integration with the host spacecraft, focusing mainly on the end-to-end communication system. A companion paper covers the details of the pointing, acquisition, and tracking (PAT) system, as well as the transmit and receive optical subassembly.¹⁴

2. PAYLOAD ARCHITECTURE AND BUILD

A high level diagram of the TBIRD architecture for the PTD-3 mission is shown in Figure 1, along with a picture of the payload prior to integration with the host CubeSat. The payload is approximately 3U in volume and weighs less than 3 kg. The optical subassembly is a bi-static design, meaning that there are separate apertures for the 200 Gbps downlink transmission (22 mm diameter) and the 2 kbps uplink beacon/data reception (23 mm diameter). The transmit beam has an as-built divergence of 450 urad full-width half-max. Both optical systems are mounted in a monolithic titanium housing that was designed to minimize transmit/receive misalignments during launch and thermal drift during operations that draw high power.




Figure 1: High-level block diagram of the TBIRD PTD-3 mission. A combination of stored data and a pseudorandom binary sequence (PRBS) are transmitted on a 200 Gbps downlink. A 2 kbps uplink signal enables error-free transmission using an automatic repeat request (ARQ) protocol. The downlink beam is precision body-pointed by the spacecraft using pointing feedback from the uplink signal.

The TBIRD communication system architecture is enabled by various COTS components and subsystems that were primarily intended for use in terrestrial applications, not space-based lasercom. The major COTS devices used include fiber telecom transceivers, NAND flash solid-state drives (SSDs), and an erbium-doped fiber amplifier (EDFA). Prior to the flight build, extensive testing was done on equivalent device units to assess their suitability for space environments and minimize the risk of issues arising during testing of the assembled payload.¹⁵ A Xilinx Ultrascale FPGA handles the various high speed interfaces between these core COTS devices and implements the high-rate data frame processing.

At the heart of the system are two 100-Gbps COTS fiber transceivers that ingest Ethernet frames and output 1550nm single mode fiber signals carrying dual-polarization QPSK waveforms for downlink transmission. The two fiber signals are wavelength-division-multiplexed and sent to the EDFA, which outputs 800mW of optical power. The transceivers are highly-integrated devices that achieve excellent performance in static fiber channels,

with receiver sensitivities around 5 photons per information bit. However, when they are used over a free-space atmospheric fading channel, fluctuations in received power result in dropped data frames and therefore an unreliable communication link. To overcome this limitation, the TBIRD system uses a custom ARQ protocol to acknowledge successfully received data over an uplink communication channel, thereby informing the space terminal which data needs to be retransmitted. The ARQ approach was validated in earlier testing over a 3-km horizontal atmospheric link using a prototype of the TBIRD payload.¹²

A single optical uplink signal in the 1.5- μ m band is used as both a communication channel for the ARQ messages and a beacon for spatial tracking. The uplink waveform uses 4-kslot/s binary pulse position modulation (2-PPM) with square pulses and no guard slots. ARQ messages are appended with a CRC and encoded with a Reed-Solomon (223,255) forward error correction code. The CRC allows the uplink receiver on the payload to reliably detect whether ARQ messages are valid, which is needed for the successful operation of the ARQ protocol. Accounting for the coding overhead, a user data rate of 1.8 kbps is available for ARQ signaling over the uplink. Uplink codewords are over two thousands bits long, so at this data rate the ARQ messages are only received by the payload's ARQ controller about once per second.

The uplink receiver is composed of a quadrant PIN photodiode and low-noise TIA, the output of which is digitized and sent to a microcontroller for spatial discrimination processing and data demodulation. The demodulated data is sent to the FPGA, which decodes the ARQ message and updates the ARQ controller accordingly. The spatial discrimination processing culminates in a two-axis pointing error signal that is provided to the bus attitude determination and control (ADCS) system at 10 Hz. With TBIRD feedback in the loop, pointing accuracy of the downlink is predicted to be about 30 μ rad RMS single-axis.¹⁴

A 2-TB memory module was built using four PC-grade COTS SSDs in parallel, each of which can read data out at 25 Gbps for an aggregate readout rate of 100 Gbps. Along with other COTS components, the SSDs were radiation-tested at the component level with gamma and proton radiation and found to be suitable for the mission. In addition to reading out at 100 Gbps, the memory module also handles writing and storing payload telemetry during operation, but this occurs at a relatively low rate. Since there is no sensor on-board that generates data at 100 Gbps, a large section of the storage contains pre-loaded data that will be used to validate the end-to-end reliable transfer of data to ground during the mission.

In the parallelized architecture, the total addressable space available is divided into ARQ data frames, each of which is around 15 MB and split over the four drives. These frames are the highest level units that are requested by the FPGA during the ARQ protocol operation, but those high level requests are divided into lower-level data blocks when interfacing with the drives. One of the key aspects of the buffer implementation was the optimization of the interface with the COTS memory controller embedded in the SSD, which was needed to achieve high read speeds while executing the random-access ARQ protocol.

The payload supports both 100-Gbps and 200-Gbps downlink modes by enabling either one or both of the 100-Gbps transceivers. In 100-Gbps mode, data can be sourced from the 2 TB buffer at 100 Gbps, or instead from a virtual buffer through the use of a pseudorandom binary sequence (PRBS). In 200-Gbps mode, the full bandwidth is achieved by reading from both sources, or by using PRBS alone.

In full operation with all components enabled, the payload draws around 100W of electrical power. Due to the size and mass constraints of the CubeSat host platform, thermal mitigation methods are limited. Instead of integrating components in a traditional layered stack, the TBIRD payload was built such that components with the highest power consumption are positioned near the exterior faces of the enclosure to achieve higher thermal radiation into space. Even so, the payload tends to self-heat after starting up in full power mode. This is acceptable for the PTD-3 mission because passes are only a few minutes in duration and occur at a very low duty cycle.

3. PAYLOAD TESTING AND RESULTS

After assembly of the payload, shock and vibration and thermal vacuum (TVAC) environmental stress tests were performed. This paper will focus on results of TVAC testing as it provides the most comprehensive test of expected payload performance on orbit. This section describes the TVAC test configuration at a high level and presents results from the end-to-end communication system tests. Tests results related to the PAT system, such as transmit/receive alignment, field of view, and, spatial tracking, are covered in the companion paper.¹⁴

3.1 Test Setup

The TVAC test setup is shown in Figure 2. The ground terminal surrogate was an engineering unit version of the payload with the addition of a low-noise pre-amp before a wavelength demultiplexer. The downlink path used a fiber test port that tapped off of the payload transmitter and passed through an all-fiber fade emulator before going to the ground receiver. The fade emulator implements a given fade time series profile using a variable optical attenuator.¹⁶ A particularly stressing fade profile was chosen for the testing, with a small coherence time ($< 1\text{ms}$) to emulate worst-case LEO-to-Ground conditions and a large scintillation index (1.0) to represent aggressive conditions at low elevation angles ($20\text{-}30^\circ$).




Figure 2: Thermal-vacuum (TVAC) test configuration. End-to-end communication tests were performed with atmospheric fading emulation on both the downlink and uplink. A free-space optical test set was used to illuminate the payload receiver with the uplink beam.

The ground terminal receives the faded signal and implements the receiver portion of the ARQ protocol. An ARQ controller determines which ARQ frames have been successfully received and creates ARQ messages to be sent over the uplink. Frames are not written to an actual storage buffer in the ground terminal surrogate, but the receiver is able to verify that the intended data transfer is ultimately successful once all of the constituent frames have been received.

The ground terminal encodes the ARQ message into uplink codewords and produces the PPM waveform by modulating an optical switch whose input is a CW laser source. The uplink fiber signal is then sent through the fade emulator, which implements a separate fade profile intended to represent the expected Ground-to-LEO fading channel. A free-space optical test set couples the fiber to a flat-top free space beam that illuminates the payload uplink receiver.

3.2 End-to-end test results

The end-to-end test was a real-time transfer of data from the payload’s memory module to the ground terminal. A section of the on-board buffer corresponding to 100,000 ARQ frames was transmitted until the entire section was received error-free by the ground terminal. In 100-Gbps mode, this number of frames corresponded to a 1.6 TB section of the 2-TB SSD buffer being delivered. In 200-Gbps mode, it corresponded to 3.4 TB of the virtual buffer being delivered. Based on the total time that it took to complete the data transfer, the average throughput in Gbps was calculated and compared to system modeling predictions.

Variations on this test were performed during hot and cold dwells in the TVAC chamber, at different average received power levels and in different modes of operation. The 100-Gbps tests used the SSDs as the data source, while the 200-Gbps tests used all PRBS. Figure 3 shows results from the testing in terms of average throughput versus average received power. The modeling and the measurements generally show good agreement, and there was very little variation throughout the testing. One exception is that the Post-TVAC performance improved slightly compared to prior testing for 200-Gbps mode. This occurred because in that test, the transmit power of the two transceivers was adjusted to optimize the receiver sensitivity (due to unbalanced EDFA gain at the two wavelengths).

At lower average power, the throughput decreases, as predicted by the model. In this regime, more frames are dropped by the receiver due to fading, and therefore the payload must spend time retransmitting those lost frames. In the absence of fading, the throughput would be the maximum 100 Gbps (or 200 Gbps), which requires roughly -42 dBm (or -39 dBm) of received power.




Figure 3: End-to-end test results. Red points correspond to 100-Gbps mode with the SSDs as the data source. Blue points correspond to 200-Gbps mode with PRBS as the data source.

3.3 Uplink and ARQ robustness results

The tests in Figure 3 were all at sufficiently high uplink power that no uplink codewords were dropped due to fading, and therefore no ARQ messages were lost. The impact of dropping a transmitted ARQ message is that the payload wastefully proceeds to retransmit certain downlink frames, thinking they were not received by the ground terminal because the acknowledgement did not make it through. A rough model of the performance impact of uplink codeword errors on the downlink transmission is that the average throughput is scaled by a factor equal to $1 - \alpha$, where α is the uplink codeword error rate (CWER). The data reliability, however, is not impacted because uplink codeword errors are detected by the CRC, and the ARQ protocol is robust to dropped ARQ messages.

To validate the robustness of the ARQ protocol and to measure the performance impact of uplink errors, the system was tested at low uplink power near the FEC threshold of the payload’s receiver. This threshold is different depending on whether the SSDs are enabled or not. When the SSDs were disabled, the uplink receiver threshold was measured to be roughly -70 dBm. When SSD were enabled and reading out data, there was electrical interference that increased the noise seen by the sensitive quadrant photodiode TIA circuit. As a result, the receive threshold measured in this regime was roughly -60 dBm.

The test results are shown in Figure 4 for the case of operating in 200-Gbps mode with the data source all PRBS (SSDs were disabled) and average received power of -34 dBm. The uplink power was adjusted such that CWERs of $\{0, 0.6\%, 43\%\}$ were measured. The plot shows terabytes transmitted and received over the course of the test, stopping when the ground terminal successfully received the 3.4 TB section of the virtual buffer that the payload attempted to transmit. In all cases, even when over 40% of the ARQ messages were dropped, the test completed successfully, validating the robustness of the ARQ protocol. The results also reveal that the system can withstand around 0.1-1% CWER on the uplink without impacting the downlink performance significantly. Relative to the CWER=0 case, the throughput factor measured for the other two cases was 96% and 73%, whereas the rough model $1 - \alpha$ gives 99% and 57%, respectively.




Figure 4: ARQ robustness test results when uplink received power is low enough to produce uplink codeword errors. Tests were performed in 200-Gbps mode with an all-PRBS data source.

4. CONCLUSION

The TBIRD PTD-3 mission on a 6U CubeSat in LEO will demonstrate a new architecture for delivering TB-class data volumes to ground. A 3U lasercom payload that leverages a variety of COTS technologies has been built and tested. End-to-end testing of the payload has shown that the system is capable of reliably transferring terabytes of data in a few minutes while operating in a space-like environment and through stressing atmospheric fading channels. The TBIRD payload has been delivered for integration with the host spacecraft which is scheduled to launch in June 2022.

REFERENCES

- [1] M. Toyoshima, “Recent trends in space laser communications for small satellites and constellations,” *IEEE Journal of Lightwave Technology* **39**(3), pp. 693–699, 2020.
- [2] B. L. Edwards, D. J. Israel, and S. K. Vithlani, “Latest changes to NASA’s laser communications relay demonstration project,” in *Free-Space Laser Communications XXX, Proc. SPIE 10524*, 2018.
- [3] F. Heine *et al.*, “Status of tesat lasercomms activities,” in *Free-Space Laser Communications XXXIII, Proc. SPIE 11678*, 2021.
- [4] “SpaceX launches first full batch of laser-equipped Starlink satellites.” <https://spaceflightnow.com/2021/09/14/spacex-launches-first-full-batch-of-laser-equipped-starlink-satellites/>. Accessed: 2022-01-10.
- [5] K. Iwamoto *et al.*, “Experimental results on in-orbit technology demonstration of SOLISS,” in *Free-Space Laser Communications XXXIII, Proc. SPIE 11678*, 2018.
- [6] C. Fuchs *et al.*, “OSIRISv1 on flying laptop: Measurement results and outlook,” in *IEEE International Conference on Space Optical Systems and Applications*, 2019.
- [7] T. S. Rose *et al.*, “Optical communications downlink from a 1.5 U CubeSat: OCSD program,” in *International Conference on Space Optics*, **11180**, p. 111800J, 2018.
- [8] B. Rödiger *et al.*, “High data-rate optical communication payload for CubeSats,” in *Laser Communication and Propagation through the Atmosphere and Oceans IX, Proc. SPIE 11506*, 2020.

- [9] D. M. Boroson *et al.*, “A new optical communication architecture for delivering extremely large volumes of data from space to ground,” *American Institute of Aeronautics and Astronautics SPACE Conference* **4658**, 2015.
- [10] B. S. Robinson *et al.*, “TeraByte InfraRed Delivery (TBIRD): A demonstration of large-volume direct-to-earth data transfer from low-earth orbit,” *Proc. SPIE* **10524**, 2018.
- [11] C. Schieler *et al.*, “NASA’s terabyte infrared delivery (TBIRD) program: Large-volume data transfer from LEO,” in *AIAA/USU Conference on Small Satellites*, 2019.
- [12] C. M. Schieler *et al.*, “Demonstration of reliable high-rate optical communication over an atmospheric link using ARQ,” in *IEEE International Conference on Space Optical Systems and Applications*, 2019.
- [13] “Pathfinder Technology Demonstrator: Demonstrating Novel CubeSat Technologies in Low-Earth Orbit.” https://www.nasa.gov/directorates/spacetech/small_spacecraft/Pathfinder_Technology_Demonstrator/. Accessed: 2022-01-10.
- [14] K. M. Riesing *et al.*, “Pointing, acquisition, and tracking for the TBIRD CubeSat mission: system design and pre-flight results,” in *Free-Space Laser Communications XXXIV, Proc. SPIE* **11993**, 2022.
- [15] N. W. Spellmeyer, “Use of commercially available fiber optic components in emerging space laser communications applications: Optimizing performance, cost, and reliability,” in *IEEE Optical Fiber Communications Conference and Exhibition (OFC)*, 2021.
- [16] D. Geisler *et al.*, “Atmospheric emulation and testing methodology for laboratory verification of FSO communications transceivers,” in *IEEE International Conference on Space Optical Systems and Applications*, IEEE, 2019.